

CURRENT-STATE CONSTRAINED FILTER BANK FOR WALD TESTING OF SPACECRAFT CONJUNCTIONS

J. Russell Carpenter⁽¹⁾ and F. Landis Markley⁽²⁾

⁽¹⁾Navigation and Mission Design Branch, Code 595, NASA Goddard Space Flight Center,
Greenbelt, MD 20771, 301-286-7526, Russell.Carpenter@nasa.gov

⁽²⁾Attitude Control Systems Engineering Branch, Code 591, NASA Goddard Space Flight Center,
Greenbelt, MD 20771, 301-286-4573, Landis.Markley@nasa.gov

Abstract: We propose a filter bank consisting of an ordinary current-state extended Kalman filter, and two similar but constrained filters: one is constrained by a null hypothesis that the miss distance between two conjuncting spacecraft is inside their combined hard body radius at the predicted time of closest approach, and one is constrained by an alternative complementary hypothesis. The unconstrained filter is the basis of an initial screening for close approaches of interest. Once the initial screening detects a possibly risky conjunction, the unconstrained filter also governs measurement editing for all three filters, and predicts the time of closest approach. The constrained filters operate only when conjunctions of interest occur. The computed likelihoods of the innovations of the two constrained filters form a ratio for a Wald sequential probability ratio test. The Wald test guides risk mitigation maneuver decisions based on explicit false alarm and missed detection criteria. Since only current-state Kalman filtering is required to compute the innovations for the likelihood ratio, the present approach does not require the mapping of probability density forward to the time of closest approach. Instead, the hard-body constraint manifold is mapped to the filter update time by applying a sigma-point transformation to a projection function. Although many projectors are available, we choose one based on Lambert-style differential correction of the current-state velocity. We have tested our method using a scenario based on the Magnetospheric Multi-Scale mission, scheduled for launch in late 2014. This mission involves formation flight in highly elliptical orbits of four spinning spacecraft equipped with antennas extending 120 meters tip-to-tip. Eccentricities range from 0.82 to 0.91, and close approaches generally occur in the vicinity of perigee, where rapid changes in geometry may occur. Testing the method using two 12,000-case Monte Carlo simulations, we found the method achieved a missed detection rate of 0.1%, and a false alarm rate of 2%.

Keywords: conjunction analysis, constrained Kalman filtering, decision theory

1. Introduction

The current state of the practice for conjunction assessment (CA) is predominantly based on attempts to explicitly compute collision probability [1], [2], [3]. In principal, such approaches require approximate solutions to the Fokker-Planck-Kolmogorov partial differential equation for mapping probability densities through time. They then require approximations to integrals of probability density to compute a collision probability. Once this estimate is in hand, they require thresholding of acceptable collision probability values, or other arbitrary factors associated with the character of the conjunction.

Reference [4] proposed the use of a Wald Sequential Probability Ratio Test [5] (WSPRT) to guide the collision avoidance decision process. The WSPRT guides risk mitigation maneuver decisions

based on explicit false alarm and missed detection criteria, which can be informed by Bayesian utility theory. Some limitations of the method proposed in Reference [4] include assumptions that the observations are statistically independent, and its reliance on a set of assumptions that reduce the complexity of the encounter. These limitations were overcome in Reference [6] which reformulated the WSPRT using a filter bank consisting of two norm-inequality-constrained epoch-state extended Kalman filters. In that approach one filter models a null hypothesis that the miss distance between two conjuncting spacecraft is inside their combined hard body radius at the predicted time of closest approach, and one is constrained by an alternative complementary hypothesis. The epoch-state filter developed for that method explicitly accounts for any process noise present in the system. Because of its epoch-state formulation however, that method still required potentially inaccurate approximations to mapping probability density forward through time.

In this work, we report our development of a filter bank that eliminates the need for an epoch-state formulation. Our current approach consists of an ordinary current-state extended Kalman filter (EKF), and two current-state constrained EKFs, one for each hypothesis in the WSPRT. The unconstrained filter is the basis of an initial screening for close approaches of interest. Once the initial screening detects a risky conjunction, the unconstrained filter predicts the time of closest approach, t_{ca} , for all three filters.¹ The constrained filters operate only when conjunctions of interest occur. The densities governing the innovations of the two constrained filters form the likelihood ratio for the WSPRT at the time of the current measurement. The unconstrained filter governs measurement editing for all three filters, avoiding any ambiguity in computing the likelihood ratio. Figure 1 provides an overview of the architecture of the proposed filter bank.

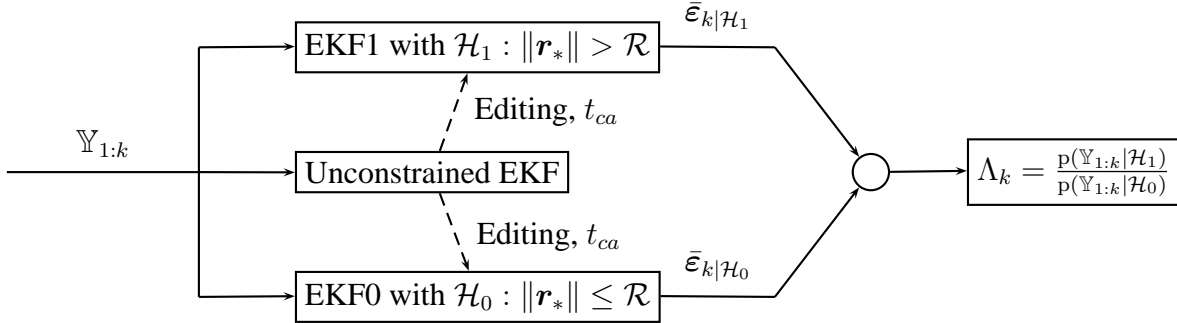


Figure 1. Proposed Current-State Filter Bank for WSPRT. The set of measurements $\mathbb{Y}_{1:k}$ is processed by two inequality-constrained EKFs, and one unconstrained EKF. The conditional densities of the innovations of the constrained filters, $\bar{\epsilon}_{k|H_i}$, are used in a likelihood ratio, Λ_k , for a WSPRT.

Since only current-state Kalman filtering is required to compute the innovations for the likelihood ratio, the present approach does not require the mapping of probability density forward to the time of closest approach. Instead, the hard-body constraint manifold is mapped to the filter update time by means of a sigma-point transformation and a projection operator. Although a multiplicity of projectors are possible, we choose one based on Lambert-style differential correction of the current-state velocity.

¹The use an unconstrained filter on which to base the initial screening and t_{ca} predictions was implicit in Reference [6]; we explicitly acknowledge its necessity here to better clarify our method.

2. Problem Description

Let $\{\mathbf{R}_i, \mathbf{V}_i\}, i = 1, 2$ represent the inertial position and velocity of two space objects. Then, the relative position and velocity between the objects in a frame fixed to the orbital velocity and orbit normal vectors of object 1 are

$$\mathbf{r} = \mathbf{R}_2 - \mathbf{R}_1 \quad \text{and} \quad \mathbf{v} = \mathbf{V}_2 - \mathbf{V}_1 - \boldsymbol{\omega}_1 \times \mathbf{r} \quad (1)$$

where $\boldsymbol{\omega}_1$ is the instantaneous orbital angular velocity of object 1. During a time interval of interest, the relative position may have one or more local minima that occur when

$$\mathbf{r}^T \mathbf{v} = 0 \quad \text{and} \quad \mathbf{r}^T \mathbf{a} + \mathbf{v}^T \mathbf{v} > 0 \quad (2)$$

where \mathbf{a} is the difference between the total inertial accelerations of the objects. Such minima are predicted as part of a CA screening process, and any minima for which some additional screening criteria are met, such as $\|\mathbf{r}\|$ less than some specification, constitute events of interest for further assessment using the method of the present work. It is common to denote as t_{ca} the time at which Eqs. 2, along with any additional screening criteria, are satisfied. In the sequel, we abbreviate this notation to t_* , and any other object with this subscript should be understood to refer to that object's value at $t_* = t_{ca}$.

Given such a conjunction event of interest, we seek to establish a decision procedure for when to recommend a collision risk mitigation maneuver. The procedure we recommend below is to maneuver based on the likelihood that measurements of the object states are consistent with the hypothesis that $\|\mathbf{r}\| \leq \mathcal{R}$, where \mathcal{R} is the combined hard-body radius of the conjuncting objects. Collecting (according to some design unimportant for the present argument) the object states at time t_k into a vector \mathbf{x}_k , we assume there exist a sequence of k measurements of these states, which we model as

$$\mathbf{y}_k = \mathbf{h}(\mathbf{x}_k) + \mathbf{v}_k \quad (3)$$

where we assume the measurement noise is a zero-mean Gaussian process with covariance R_k , which we will indicate by expressing its probability density as $p(\mathbf{v}_k) = \mathcal{N}(\mathbf{v}_k, 0, R_k)$, where for $\mathbf{x} \in \mathbb{R}^n$, $\boldsymbol{\mu} = \mathbb{E}[\mathbf{x}]$, and $P = \mathbb{E}[(\mathbf{x} - \boldsymbol{\mu})(\mathbf{x} - \boldsymbol{\mu})^T]$,

$$\mathcal{N}(\mathbf{x}, \boldsymbol{\mu}, P) = (2\pi)^{-n/2} |P|^{-1/2} e^{-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T P^{-1}(\mathbf{x} - \boldsymbol{\mu})} \quad (4)$$

3. Wald Sequential Probability Ratio Test

As in References [4] and [6], we employ the WSPRT for our decision procedure. The WSPRT uses a ratio of the joint probability densities of the set of k measurements of the spacecraft, $\mathbb{Y}_{1:k} = \{\mathbf{y}_1, \dots, \mathbf{y}_k\}$, under the alternative hypothesis, \mathcal{H}_1 that the conjunction is safe, and the null hypothesis, \mathcal{H}_0 , that the conjunction is unsafe:

$$\Lambda_k = \frac{p(\mathbb{Y}_{1:k} | \mathcal{H}_1)}{p(\mathbb{Y}_{1:k} | \mathcal{H}_0)} = \frac{p(\mathbb{Y}_{1:k} | \|\mathbf{r}_*\| > \mathcal{R})}{p(\mathbb{Y}_{1:k} | \|\mathbf{r}_*\| \leq \mathcal{R})} \quad (5)$$

In a Wald test, one compares Λ_k to decision limits A and B such that whenever $B < \Lambda_k < A$ one should, if possible, seek another observation. If $\Lambda_k \leq B$, then one should accept the null hypothesis, and in the present case, we would recommend a collision avoidance maneuver. If $\Lambda_k \geq A$, then one should accept the alternative hypothesis, and hence we would dismiss the conjunction alarm². Wald's explanations for the thresholds A and B are that we will accept the alternative hypothesis if it is A times more likely than the null, and accept the null hypothesis if it is $1/B$ times more likely than the alternative. Wald shows that such a procedure will terminate with probability one, and that

$$A \leq \frac{1 - P_{fa}}{P_{md}} \quad \text{and} \quad B \geq \frac{P_{fa}}{1 - P_{md}} \quad (6)$$

where P_{fa} is the allowable false alarm probability, and P_{md} is the allowable missed detection probability.

3.1. Joint Density of the Measurements

The joint density of the measurements, unconstrained by either hypothesis, can be written in a sequential form as

$$p(\mathbb{Y}_{1:k}) = p(\mathbf{y}_k | \mathbb{Y}_{1:k-1}) p(\mathbb{Y}_{1:k-1}) \quad (7)$$

As standard texts, such as Brown and Hwang [7], show, the conditional density of the k th measurement conditioned on the past measurement sequence can be written in terms of the innovations of a sequential estimator, such as the Kalman filter. If the noise inputs to the estimator are zero-mean and Gaussian, then

$$p(\mathbf{y}_k | \mathbb{Y}_{1:k-1}) = N(\mathbf{y}_k, \mathbf{h}(\bar{\mathbf{x}}_k), \bar{H}_k \bar{P}_k \bar{H}_k^T + R_k) \quad (8)$$

where $\bar{\mathbf{x}}_k$ is the filter's estimate at t_k just prior to incorporating the measurement \mathbf{y}_k , $\mathbf{y}_k - \mathbf{h}(\bar{\mathbf{x}}_k) \equiv \bar{\mathbf{e}}_k$ is the k th filter innovation, $\bar{H}_k \bar{P}_k \bar{H}_k^T + R_k \equiv \bar{W}_k$ is the innovations covariance, \bar{P}_k is the filter's covariance corresponding to errors in $\bar{\mathbf{x}}_k$, and $\bar{H}_k = \partial \mathbf{h}_k / \partial \mathbf{x}_k |_{\bar{\mathbf{x}}_k}$. Similarly, we may distinguish the filter's k th residual as $\hat{\mathbf{e}}_k \equiv \mathbf{y}_k - \mathbf{h}(\hat{\mathbf{x}}_k)$, with $\hat{H}_k \hat{P}_k \hat{H}_k^T + R_k \equiv \hat{W}_k$ as the residual covariance, \hat{P}_k as the filter's covariance corresponding to errors in $\hat{\mathbf{x}}_k$, and $\hat{H}_k = \partial \mathbf{h}_k / \partial \mathbf{x}_k |_{\hat{\mathbf{x}}_k}$.

For the likelihood ratio, we need the joint density of the measurements constrained by hypothesis \mathcal{H}_i , which we write similarly as

$$p(\mathbb{Y}_{1:k} | \mathcal{H}_i) = p(\mathbf{y}_k | \mathbb{Y}_{1:k-1}, \mathcal{H}_i) p(\mathbb{Y}_{1:k-1} | \mathcal{H}_i) \quad (9)$$

and we express $p(\mathbf{y}_k | \mathbb{Y}_{1:k-1}, \mathcal{H}_i)$ in terms of the innovations of a filter constrained by \mathcal{H}_i as

$$p(\mathbf{y}_k | \mathbb{Y}_{1:k-1}, \mathcal{H}_i) = N(\mathbf{y}_k, \mathbf{h}(\bar{\mathbf{x}}_{k|\mathcal{H}_i}), \bar{H}_{k|\mathcal{H}_i} \bar{P}_{k|\mathcal{H}_i} \bar{H}_{k|\mathcal{H}_i}^T + R_k) \quad (10)$$

with an obvious extension of the notation above.

²In the present case, there may be minimal penalty in waiting until all possible measurements have been collected. If the test is still indeterminate at that time, it may be prudent to maneuver, although this would imply an increased false alarm rate.

3.2. Constraint Mapping

To constrain the epoch-state filters in Reference [6] we extended the method for equality-constrained estimation of Zanetti et al. [8] to the case of inequality constraints through the use of a slack variable, but this approach does not appear feasible for the current-state filter formulation. Instead, we follow the guidance of Julier and LaViola [9], in which a projection is applied to the unconstrained mean and covariance using a sigma-point transformation. First, we perturb each constrained filter's state estimate, $\hat{\mathbf{x}}$, using the columns of the Cholesky factorization of its error covariance to generate a set of sigma points, $\mathcal{X}_{k|\mathcal{H}_i}$. In the present case, the (inequality) constraint function is given by

$$c(\mathbf{x}) = \|\phi_*^r(\mathbf{x})\| \leq \mathcal{R} \quad (11)$$

where $\phi_*^r(\mathbf{x})$ is the prediction of the relative position component of \mathbf{x} to t_{ca} . Reference [9] concerns equality-constrained estimation, while we need inequality-constrained estimates. As in Reference [6], if $c(\mathcal{X}_{k|\mathcal{H}_0}^{(j)}) \geq \mathcal{R}$ for any sigma point arising from the null hypothesis filter, then the sigma point is projected onto the constraint boundary, and vice versa for the alternative hypothesis filter. Julier and LaViola's method for constraining the estimate allows for any projection function, $\mathbf{p}(\mathbf{x})$, that satisfies the constraint, i.e.

$$c(\mathbf{p}(\mathbf{x})) = \mathcal{R}. \quad (12)$$

We therefore evaluate each sigma point, $\mathcal{X}_{k|\mathcal{H}_i}^{(j)}$, compare it to the constraint, and apply the projector:

$$\mathcal{P}_{k|\mathcal{H}_i}^{(j)} = \begin{cases} \mathbf{p}(\mathcal{X}_{k|\mathcal{H}_i}^{(j)}) & \text{if } c(\mathcal{X}_{k|\mathcal{H}_i}^{(j)}) \in \mathcal{H}_i \\ \mathcal{X}_{k|\mathcal{H}_i}^{(j)} & \text{otherwise} \end{cases} \quad (13)$$

Applying the projector, $\mathbf{p}(\mathbf{x})$, is not the innovation of Julier and LaViola's method, but rather only the first step of a two-step procedure³. They argue the need for a second step by pointing out that projection does not necessarily force the mean of the resulting constrained density to satisfy the constraint. Since the mean is the property of the distribution estimated by a filter, Ref. [9] suggests that the entire distribution should be translated so that the mean will satisfy the constraint. This translation, which also affects the covariance, constitutes the second step of Julier and LaViola's approach. In the context of an inequality constraint, we would only perform this step if the mean fails to satisfy the inequality. It is not obvious to us that the second step of Julier and LaViola's method is necessary or even desirable in the present context. By construction, the projected sigma points, $\mathcal{P}_{k|\mathcal{H}_i}^{(j)}$, will already satisfy our inequality constraint and hence are our best available representation of the distribution. Translating them in order for the mean to satisfy the constraint could therefore degrade the overall quality of our representation of the distribution.

In the present case, we can in principle choose a projector by solving Lambert's problem for the velocity of one spacecraft at time t_k that will adjust the length of the relative position at the t_{ca} so that it lies on the constraint boundary, as necessary to satisfy the inequality constraint. For two-body motion, the solution may be expressed using classical f and g functions formulated using universal

³The nonlinear projection method was suggested in Ref. [10], but Julier and LaViola appear to be the first to suggest applying the projector using a sigma-point transformation.

variables, viz.

$$\mathbf{V}_1(t_k) = \frac{\mathbf{R}_1^*(t_*) - f\mathbf{R}_1(t_k)}{g} \quad (14)$$

where $\mathbf{R}_1^*(t_*)$ is the targeted position of one spacecraft such that the relative position at t_* is equal to the combined hard body radius:

$$\mathbf{R}_1^*(t_*) = \mathbf{R}_2(t_*) - \mathcal{R} \frac{\mathbf{r}_*}{\|\mathbf{r}_*\|} \quad (15)$$

Hence, if we now specify the contents of the state vector as $\mathbf{x} = [\mathbf{R}_1^T, \mathbf{V}_1^T, \mathbf{R}_2^T, \mathbf{V}_2^T]^T$, the projector is given by

$$\mathbf{p}(\mathbf{x}_k) = \begin{bmatrix} \mathbf{R}_1(t_k) \\ \frac{\mathbf{R}_1^*(t_*) - f\mathbf{R}_1(t_k)}{g} \\ \mathbf{R}_2(t_k) \\ \mathbf{V}_2(t_k) \end{bmatrix} \quad (16)$$

for those \mathbf{x}_k that do not satisfy the appropriate inequality constraint for the current hypothesis. Otherwise, the projector leaves states that already satisfy the constraint alone. Figure 2 schematically illustrates the operation of the projector function. Unlike Ref. [9], the present case has an inequality constraint. As the figure shows, sigma points that already satisfy the constraint are not projected.

It might be argued that projecting the relative position onto the constraint surface is an arbitrary choice. To justify it, we appeal to the following argument. Since we are using the Euclidean two-norm, the minimum mean-square error of the inequality-constrained estimate is given by

$$\begin{aligned} \text{MMSE} &= \arg \min_{\mathbf{x}' \in \mathcal{H}_i} \mathbb{E} [\|\mathbf{x}' - \mathbf{x}\|^2] = \arg \min_{\mathbf{x}' \in \mathcal{H}_i} \mathbb{E} [\|\mathbf{x}' - \hat{\mathbf{x}} + \hat{\mathbf{x}} - \mathbf{x}\|^2] \\ &= \arg \min_{\mathbf{x}' \in \mathcal{H}_i} \left\{ \mathbb{E} [\|\mathbf{x}' - \hat{\mathbf{x}}\|^2] + \mathbb{E} [\|\hat{\mathbf{x}} - \mathbf{x}\|^2] - 2 \mathbb{E} [(\mathbf{x}' - \hat{\mathbf{x}})^T (\hat{\mathbf{x}} - \mathbf{x})] \right\} \end{aligned} \quad (17)$$

and since $\hat{\mathbf{x}} = \mathbb{E} [\mathbf{x} | \mathbb{Y}_{1:k}]$, then

$$\text{MMSE} = \arg \min_{\mathbf{x}' \in \mathcal{H}_i} \mathbb{E} [\|\mathbf{x}' - \hat{\mathbf{x}}\|^2] \quad (18)$$

and hence the MMSE is achieved by letting \mathbf{x}' be as close as possible to $\hat{\mathbf{x}}$ in the sense of the chosen norm, subject to the constraint that $\mathbf{x}' \in \mathcal{H}_i$. If $\hat{\mathbf{x}} \in \mathcal{H}_i$, then we merely set $\mathbf{x}' = \hat{\mathbf{x}}$. Since the constraint is a boundary for \mathcal{H}_i , then if $\hat{\mathbf{x}} \notin \mathcal{H}_i$, then we choose the point on the constraint closest to $\hat{\mathbf{x}}$ in the sense of the chosen norm.

In practice, we note that a Lambert solution such as we have just described will not project estimates accurately enough under real-world perturbations for the constraints to be satisfied with adequate precision, so some kind of differential correction scheme will be required. If the conjunction is more than one orbit period in the future, some kind of two-point boundary value problem solver that breaks the prediction into segments, such as a collocation method, is likely to be necessary for the general CA case. Such methods are readily available in common off-the-shelf software, and are beyond the scope of this paper.

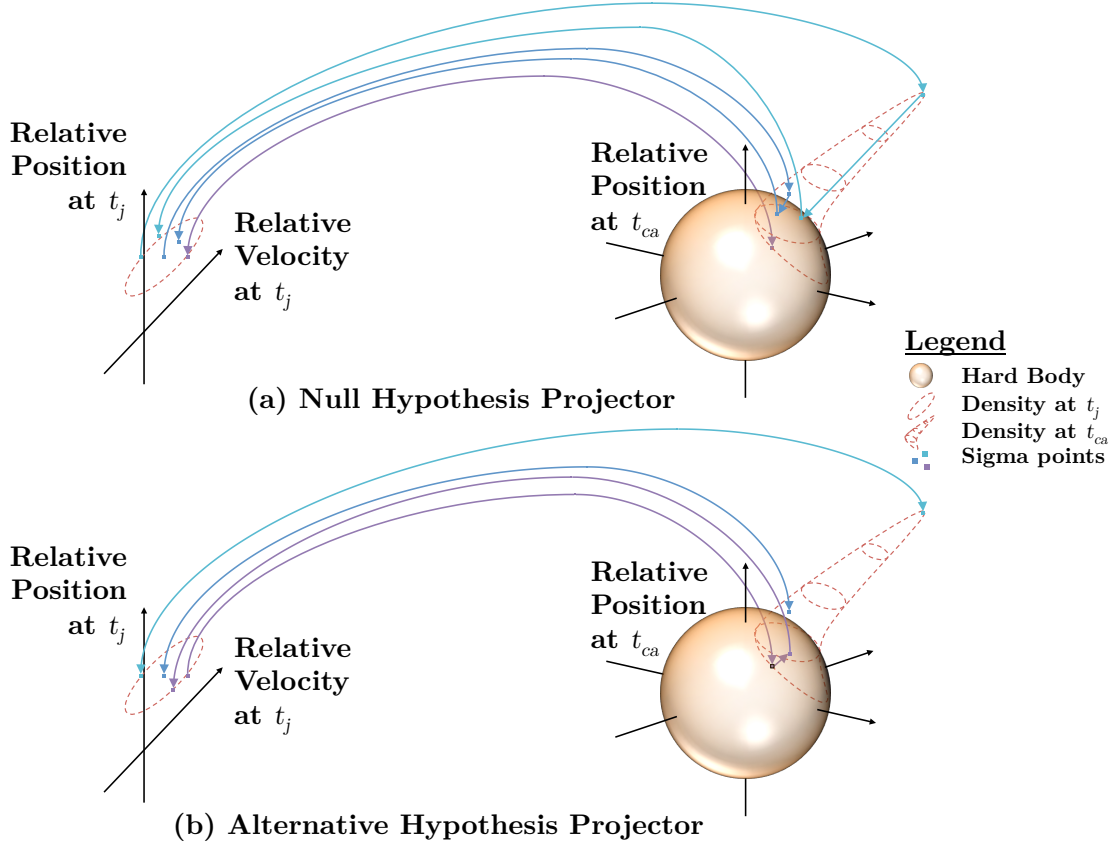


Figure 2. Schematic of Projector Mapping: (a) Null Hypothesis Projector; (b) Alternative Hypothesis Projector. Paths with a terminal arrowhead indicate how sigma points that violate the constraint are projected. Paths with arrows at each end illustrate that sigma points already satisfying the constraint are not projected. Note that $\mathbf{r} \in \mathbb{R}^3$ and $\mathbf{v} \in \mathbb{R}^3$ have been projected onto \mathbb{R}^1 at t_j , that not all sigma points are shown, and that only the velocities of sigma points at t_j are altered by the projection.

3.3. Summary of the Problem Solution

Herein, we summarize our application of the WSPRT to the CA problem.

Screening For each potential object that might be involved in a conjunction, perform orbit determination using an unconstrained estimator. For the present work, we use the extended Kalman filter. If there are no common measurements (such as might arise from inter-spacecraft tracking), no common process noise, and no common initial condition errors, one may use a separate estimator for each spacecraft. Based on some specified screening criteria, determine any pairs of objects that should be subjected to more detailed CA. The subject of appropriate screening criteria is beyond the scope of the present work, but might consist of one or more metrics, such as predicted minimum range or Mahalanobis distance, over a given planning horizon. To economize our notation, we

collect the positions and velocities estimated by the unconstrained orbit determination process(es) into a single state estimate of a notionally single extended Kalman filter. Any additional bias states, if present, would be appended to this estimate, but as they do not contribute directly to the elucidation of our algorithm, we do not consider them here. Thus, the unconstrained estimate will be given as follows. For each measurement \mathbf{y}_k , predict the estimate and its associated error covariance from the previous measurement time t_{k-1} to the current measurement time t_k :

$$\bar{\mathbf{x}}_k = \phi_k(\hat{\mathbf{x}}_{k-1}) \quad \text{and} \quad \bar{P}_k = \Phi_{k,k-1} \hat{P}_{k-1} \Phi_{k,k-1}^T + Q_k \quad (19)$$

where $\phi_k(\hat{\mathbf{x}}_{k-1})$ represents the solution of the differential equations of motion for the state over the interval $[t_k, t_{k-1}]$, $\Phi_{k,k-1}$ represents the solution of the associated variational equations, and Q_k is the covariance of any process noise input into the system over the same interval. Then, if the Mahalanobis distance from the observed measurement to the predicted measurement is within a specified editing threshold, α , update the state and covariance using Kalman's gain, $K_k = \bar{P}_k \bar{H}_k^T \bar{W}_k^{-1}$:

$$\text{if } \bar{\mathbf{e}}_k^T \bar{W}_k^{-1} \bar{\mathbf{e}}_k \leq \alpha: \begin{cases} \hat{\mathbf{x}}_k = \bar{\mathbf{x}}_k + K_k \bar{\mathbf{e}}_k \\ \hat{P}_k = (I - K_k \bar{H}_k) \bar{P}_k (I - K_k \bar{H}_k)^T + K_k R_k K_k^T \end{cases} \quad (20)$$

$$\text{otherwise:} \begin{cases} \hat{\mathbf{x}}_k = \bar{\mathbf{x}}_k \\ \hat{P}_k = \bar{P}_k \end{cases} \quad (21)$$

If there is no previous measurement, the state and covariance are initialized using prior information $\hat{\mathbf{x}}_0$ and \hat{P}_0 . After processing each measurement, predict the estimate into the future over the planning horizon, and evaluate it using the specified screening criteria, e.g. Mahalanobis distance. If the screening test indicates the need for further analysis, begin the WSPRT. Otherwise, continue screening.

Constrained Filtering To begin the WSPRT, initialize the constrained filters using the unconstrained filter state and covariance: $\hat{\mathbf{x}}_{k|\mathcal{H}_i} = \hat{\mathbf{x}}_k$ and $\hat{P}_{k|\mathcal{H}_i} = \hat{P}_k$; preferably, the unconstrained filter will have reached a “converged” status before reaching this step. Note that equating the unconstrained and constrained filters should only occur when starting up the unconstrained filters; for subsequent measurement processing, the unconstrained filters should perform their own separate time and measurement updates as follows:

$$\bar{\mathbf{x}}_{k|\mathcal{H}_i} = \phi_k(\hat{\mathbf{x}}_{k-1|\mathcal{H}_i}) \quad \text{and} \quad \bar{P}_{k|\mathcal{H}_i} = \Phi_{k,k-1|\mathcal{H}_i} \hat{P}_{k-1|\mathcal{H}_i} \Phi_{k,k-1|\mathcal{H}_i}^T + Q_k \quad (22)$$

$$\text{if } \bar{\mathbf{e}}_k^T \bar{W}_k^{-1} \bar{\mathbf{e}}_k \leq \alpha: \begin{cases} \hat{\mathbf{x}}_{k|\mathcal{H}_i} = \bar{\mathbf{x}}_{k|\mathcal{H}_i} + K_{k|\mathcal{H}_i} \bar{\mathbf{e}}_{k|\mathcal{H}_i} \\ \hat{P}_{k|\mathcal{H}_i} = (I - K_{k|\mathcal{H}_i} \bar{H}_{k|\mathcal{H}_i}) \bar{P}_{k|\mathcal{H}_i} (I - K_{k|\mathcal{H}_i} \bar{H}_{k|\mathcal{H}_i})^T + K_{k|\mathcal{H}_i} R_k K_{k|\mathcal{H}_i}^T \end{cases} \quad (23)$$

$$\text{otherwise:} \begin{cases} \hat{\mathbf{x}}_{k|\mathcal{H}_i} = \bar{\mathbf{x}}_{k|\mathcal{H}_i} \\ \hat{P}_{k|\mathcal{H}_i} = \bar{P}_{k|\mathcal{H}_i} \end{cases} \quad (24)$$

Note also that editing for the constrained filters is controlled by the editing decisions of the unconstrained filter, and that all three filters share the same process and measurement noise covariances, but possibly differ in their state transition and measurement Jacobians.

After each measurement update, the constrained filters apply the constraint as described in Section 3.2. to the posterior state estimate. We then merge the constrained sigma points to form the “step-1” constrained mean:

$$\hat{\mathbf{x}}_{k|\mathcal{H}_i}^{(1)} = \frac{h^2 - n}{h^2} \mathcal{P}_{k|\mathcal{H}_i}^{(1)} + \frac{1}{2h^2} \sum_{j=2}^{2n+1} \mathcal{P}_{k|\mathcal{H}_i}^{(j)} \quad (25)$$

where h is the interval length or step size of the sigma point transformation, and n is the dimension of the state. To form the associated covariance, we use the divided-difference matrices of Nørgaard, et al. [11, 12], each column j of which is given by⁴

$$\tilde{D}_{\Delta x}^{(1)} \mathbf{p}(\hat{\mathbf{x}}_{k|\mathcal{H}_i})_j = \frac{1}{2h} \left[\mathcal{P}_{k|\mathcal{H}_i}^{(j+1)} - \mathcal{P}_{k|\mathcal{H}_i}^{(j+1+n)} \right] \quad (26)$$

$$\tilde{D}_{\Delta x}^{(2)} \mathbf{p}(\hat{\mathbf{x}}_{k|\mathcal{H}_i})_j = \frac{\sqrt{h^2 - 1}}{2h^2} \left[\mathcal{P}_{k|\mathcal{H}_i}^{(j+1)} + \mathcal{P}_{k|\mathcal{H}_i}^{(j+1+n)} - 2\mathcal{P}_{k|\mathcal{H}_i}^{(1)} \right] \quad (27)$$

The application of the nonlinear constraint to the covariance is then approximated (for step 1) by

$$\hat{P}_{k|\mathcal{H}_i}^{(1)} = \left[\tilde{D}_{\Delta x}^{(1)} \mathbf{p}(\hat{\mathbf{x}}_{k|\mathcal{H}_i}), \tilde{D}_{\Delta x}^{(2)} \mathbf{p}(\hat{\mathbf{x}}_{k|\mathcal{H}_i}) \right] \left[\tilde{D}_{\Delta x}^{(1)} \mathbf{p}(\hat{\mathbf{x}}_{k|\mathcal{H}_i}), \tilde{D}_{\Delta x}^{(2)} \mathbf{p}(\hat{\mathbf{x}}_{k|\mathcal{H}_i}) \right]^T \quad (28)$$

Finally, if we are performing Julier and LaViola’s step 2, we apply the projector to the step-1 mean and adjust its covariance, if necessary:

$$\text{if } c(\hat{\mathbf{x}}_{k|\mathcal{H}_i}^{(1)}) \in \mathcal{H}_i: \begin{cases} \hat{\mathbf{x}}_{k|\mathcal{H}_i}^{(2)} = \mathbf{p}(\hat{\mathbf{x}}_{k|\mathcal{H}_i}^{(1)}) \\ \hat{P}_{k|\mathcal{H}_i}^{(2)} = \hat{P}_{k|\mathcal{H}_i}^{(1)} + \left(\hat{\mathbf{x}}_{k|\mathcal{H}_i}^{(2)} - \hat{\mathbf{x}}_{k|\mathcal{H}_i}^{(1)} \right) \left(\hat{\mathbf{x}}_{k|\mathcal{H}_i}^{(2)} - \hat{\mathbf{x}}_{k|\mathcal{H}_i}^{(1)} \right)^T \end{cases} \quad (29)$$

$$\text{otherwise:} \begin{cases} \hat{\mathbf{x}}_{k|\mathcal{H}_i}^{(2)} = \hat{\mathbf{x}}_{k|\mathcal{H}_i}^{(1)} \\ \hat{P}_{k|\mathcal{H}_i}^{(2)} = \hat{P}_{k|\mathcal{H}_i}^{(1)} \end{cases} \quad (30)$$

Likelihood Ratio Finally, we iteratively compute the likelihood ratio for the WSPRT from the innovations of the constrained filters (in practice we have found that it is generally numerically superior to compute the log-likelihood ratio):

$$\Lambda_k = \frac{\mathcal{N} \left(\mathbf{y}_k, \mathbf{h}(\bar{\mathbf{x}}_{k|\mathcal{H}_1}), \bar{H}_{k|\mathcal{H}_1} \bar{P}_{k|\mathcal{H}_1} \bar{H}_{k|\mathcal{H}_1}^T + R_k \right)}{\mathcal{N} \left(\mathbf{y}_k, \mathbf{h}(\bar{\mathbf{x}}_{k|\mathcal{H}_0}), \bar{H}_{k|\mathcal{H}_0} \bar{P}_{k|\mathcal{H}_0} \bar{H}_{k|\mathcal{H}_0}^T + R_k \right)} \Lambda_{k-1} \quad (31)$$

and compare Λ to Wald’s limits A and B , computed using desired missed detection and false alarm rates. If $\Lambda > A$ we dismiss the conjunction and discontinue the constrained filter bank. If $\Lambda \leq B$ we recommend a conjunction risk mitigation maneuver.

⁴The explicit form of these equations assumes indexing the sigma points as in Refs. [11, 12].

4. Test Data

Spacecraft collisions are rare events. It would be difficult to comprehensively test the present algorithm using realistic spacecraft conjunctions, so we construct artificial conjunctions that are arguably more stressful, and certainly more comprehensive, than actual mission data could provide.

We begin by choosing a test distribution for the miss distances at a given t_{ca} . To make the present explanation definitive, we select a Maxwell distribution, which we can relate to a χ^2 distribution:

$$\chi_\nu^2(z) = [2^{\nu/2}\Gamma(\nu/2)]^{-1} z^{(\nu/2-1)} e^{-z/2} \quad (32)$$

We want to scale this distribution, with $\nu = 3$, such that the square roots of half of all samples from it (z) lie inside the combined hard body surface, and half lie outside. To achieve this, we find the value z_* that satisfies

$$0.50 = \Pr(z \leq z_*) = \int_0^{z_*} \chi_3^2(z) dz \quad (33)$$

and then we scale $\chi_3^2(z)$ samples by $\mathcal{R}/\sqrt{z_*}$.

We have found it helpful to further categorize the samples into quartiles which we denote as “clear hit,” “near hit,” “near miss,” and “clear miss.” We then distribute these miss distances uniformly over \mathbb{R}_3 using “igloo” sampling [13]. Figure 3 illustrates the resulting distribution of relative position vectors, color-coded by the quartile category. Two quadrants of data from the foreground of the figure have been excluded so as to provide a cut-away view into the interior, and hence better illustrate the distribution of data among the quartiles.

Relative velocities associated with each relative position vector sampled in this fashion need to satisfy the conditions for minimum, given by Eqs. 2. At the time of close approach the relative acceleration may be reasonably approximated as a linear function of relative position and velocity,

$$\mathbf{a} \approx C\mathbf{r} + D\mathbf{v} \quad (34)$$

where the matrix C contains the centripetal and gravity gradient terms, and D the Coriolis term, so that we can replace the second of Eqs. 2 by

$$\mathbf{r}^T C \mathbf{r} + \mathbf{r}^T D \mathbf{v} + \mathbf{v}^T \mathbf{v} \geq 0 \quad (35)$$

We now choose $\mathbf{v} = \mathbf{v}_N = v\mathbf{u}_N$ where v is a velocity magnitude typical of conjunctions of interest, and \mathbf{u}_N is a unit vector that spans the null space of the matrix

$$F = \begin{bmatrix} \mathbf{r}^T \\ \mathbf{r}^T D \end{bmatrix} \quad (36)$$

Such \mathbf{v}_N immediately satisfy the first of Eqs. 2, and we can make them satisfy the second as follows. With $\mathbf{v} = \mathbf{v}_N$, Eq. 35 becomes

$$\mathbf{r}^T C \mathbf{r} + \mathbf{v}_N^T \mathbf{v}_N \geq 0 \quad (37)$$




Figure 3. Cut-away view of the distribution of miss vectors. Half of the cases have miss distances of less than the combined hard body radius $\mathcal{R} = 120$ m, and half exceed this. One-quarter of the cases have miss distances of less than 86 m, and one-quarter exceed 158 m.

Since $\mathbf{v}_N^T \mathbf{v}_N \geq 0$ then Eq. 37 will be satisfied if we exclude any \mathbf{v}_N that fail to satisfy

$$\|\mathbf{v}_N\|^2 \geq \|C\|_2 \|\mathbf{r}\|^2 \quad (38)$$

Since $\|C\|_2 \sim \omega^2$, where ω is the instantaneous orbital rate, this condition is usually satisfied for close approach distances of interest to the present work, and we rarely need to exclude a \mathbf{v}_N .

Next, we introduce uncertainty into the t_{ca} by drawing perturbations to its nominal value from a given distribution. Figure 4 illustrates the use of Gaussian draws to perturb the t_{ca} , introducing uncertainty on the order of a minute.

5. Results

Using the procedures in the previous section, we can generate large sets of stressing conjunction test data for any orbit and conjunction scenario of interest. In the current section, we present results using such a data set based on the Magnetospheric Multi-Scale (MMS) mission, which NASA plans to launch in late 2014.

5.1. Example Description

Figure 5 depicts the MMS orbit during Phase 1 of the mission, an approximately 24-hour orbit with apogee radius of approximately 12 Earth radii (R_E), and perigee radius of approximately 1.2 R_E . MMS consists of four nearly identical spacecraft spinning at 3 RPM, which must maintain an




Figure 4. Histogram of empirical distribution of miss distance and t_{ca} .

approximately tetrahedral formation in the vicinity of apogee to make electromagnetic measurements of the interaction between the earth's and the sun's magnetic fields. Each MMS spacecraft has radial wire antennas approximately 60 m long. Hence, a combined hard-body sphere for two MMS spacecraft has a radius of 120 m.

In order to achieve a nearly regular tetrahedron near apogee, the MMS formation tends to need at least one pair of relative orbits that closely approach one another in the vicinity of perigee, often near the semi-latus recta crossings at true anomalies of 270 deg and 90 deg. The mission will require maneuvers approximately every two weeks in order to maintain formation quality and avoid close approaches. When MMS performs such maneuvers, it loses a full orbit's worth of science. Furthermore, as with most missions, fuel margins are tight, and maneuver operations are costly in terms of ground activities as well. Hence, MMS has a strong need to limit unnecessary conjunction avoidance maneuvers. At the same time, MMS has a total mission cost on the order of \$1B, and its orbit transits important regions of operation for other missions, including GEO, the orbits of GNSS satellites, and the higher regions of LEO where for example a large number of sun-synchronous missions operate. Furthermore, all NASA missions have a duty to avoid polluting the space environment with collision debris. Hence, the mission has a strong need to ensure that the probability of a collision is as low as practicable. Considerations such as these have led the mission to establish requirements that limit the false alarm rate for conjunction avoidance to 1/20, and the missed detection rate to 1/1000. The corresponding WSPRT limits for MMS are therefore $A = 950$ and $B = .05005$.

Although the method of this paper might in principle be used for avoiding conjunctions between MMS and other spacecraft, as well as for conjunctions within the MMS formation, for the present example, we model two MMS spacecraft that have a conjunction near a true anomaly of 270 deg. This point appears in Figure 5 as a blue asterisk. Figure 5 also depicts the portion of MMS' orbit

where we expect its GPS receiver to acquire four or more GPS signals, and hence be able to produce point solutions⁵. We model these point solutions as bias-free with isotropic white noise of 1 m per component. The data arc is two hours long, and measurements occur once per minute. We model the full nonlinear two-body dynamics for each spacecraft in simulating these measurements. The EKF's that process these measurements also model the full nonlinear two-body dynamics of each spacecraft, but the filters' dynamics model includes a small, fictitious acceleration process noise, with an isotropic intensity of $1 \times 10^{-9} \text{ m/s}^{5/2}$ per axis. The filters' measurement model is identical to the simulation model. Note that we have intentionally matched the filter model to the simulation model so as to remove any ambiguity concerning the performance of the WSPRT. We have chosen a simple but fairly nonlinear model so that we can run large numbers of cases with acceptable computational time. In the sequel, we present a detailed look at four of these cases, and then summarize the outcome of 24,000 simulations.

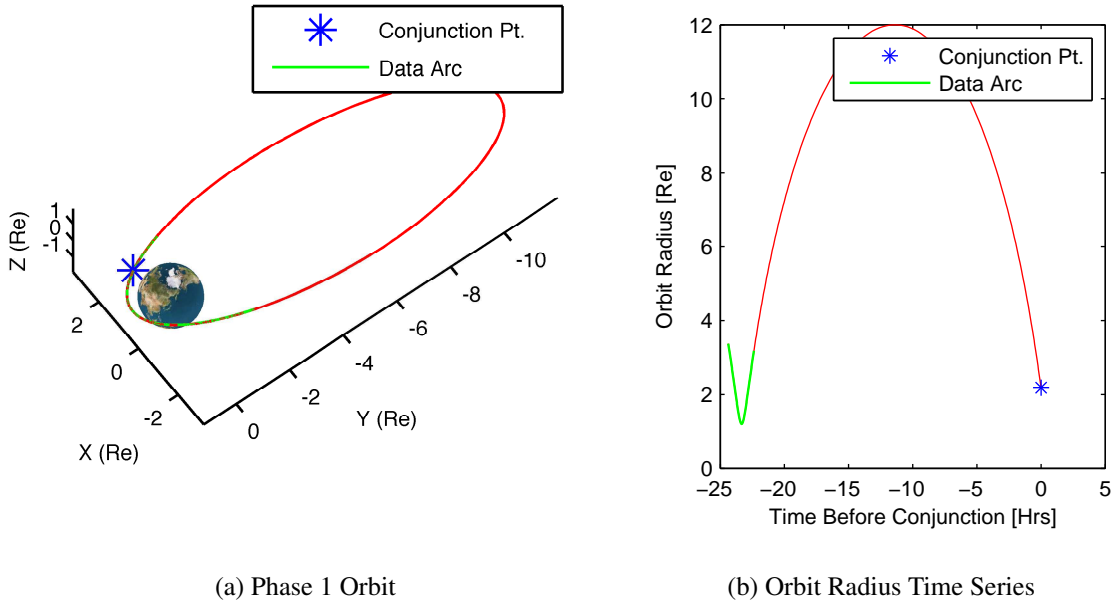


Figure 5. MMS Conjunction Example

5.2. Example Results, Part 1

To illustrate the WSPRT, we simulated four cases, one each of the clear hit, near hit, near miss, and clear miss category, which we have described above. Figures 6 to 9 summarize these results. Each figure consists of four quadrants. The upper left quadrant depicts the 120 m hard-body sphere as a gray wire mesh. Markers indicated by the legend show the relative position vector that each of the three EKF's estimates at the end of the data arc, along with the true relative position. The lower left

⁵The actual MMS GPS receiver will acquire and process, using an onboard EKF, pseudo-range data throughout all of the Phase 1 orbit, and much of the higher-apogee Phase 2 orbit so that state updates will be available over a much larger fraction of the orbit than this example considers.

and upper right quadrants show, as a solid colored line, the three components of the relative position estimation error from the alternative (\mathcal{H}_1) and null hypothesis (\mathcal{H}_0) filters, respectively. Dashed lines show the filters' 3σ formal error envelope. For reference, gray solid and dashed lines show the unconstrained filter's estimate and formal error envelope. The lower right quadrant shows the base ten logarithm of the likelihood ratio, along with Wald's limits for the alarm, B , and the dismissal, A .

In Figure 6, which is a clear hit case, we see that the null hypothesis filter closely tracks the unconstrained filter leading to a rapid and definitive conclusion to raise the alarm and recommend a conjunction avoidance maneuver. Figure 7, which is a near hit case, produces similar results. In Figure 8, which is a near miss case, neither of the constrained filters predicts as accurately as the unconstrained filter, but the alternative hypothesis filter is slightly more consistent with its own formal errors than the null hypothesis filter, leading to a dismissal. Figure 9, the clear miss case, shows a close agreement between the alternative hypothesis and unconstrained filters, and divergence of the null hypothesis filter, once again leading to a dismissal.

5.3. Example Results, Part 2

Next, we shortened the simulation timeline to consider an “emergency” situation such as might occur after a faulty formation maintenance maneuver. In this case, the faulty maneuver has occurred prior to the previous apogee, creating an unintentional close approach on the approach to the subsequent perigee. We give the WSPRT access to only 20 minutes of measurement data that begin 40 minutes prior to the conjunction, which as previously occurs at a true anomaly of 270 deg. Although this is an unrealistically short timeline for MMS to actually perform an emergency maneuver, we want to show that our WSPRT can perform adequately in even such an unrealistically stressing case. We also used this shorter scenario to examine the difference in performance between using both steps of Julier and LaViola's constraint application approach, versus only performing the projection step. We ran 12,000 Monte Carlo trials for this scenario for each of these two cases.

Table 1. Large Ensemble (12,000 cases each) Stress Case Results

Case	Missed Detection	False Alarm	No Decision	Indecision	Eff. False Alarm
Both Steps	0.167%	2.17%	10.0%	0.7%	7.35%
Step 1 Only	0.117%	1.82%	9.63%	0.5%	7.59%

In the first row of Table 1 are the results of using both steps (projection and translation) of the constraint application, while the second row contains the results of using the projection step only. The first two columns of the table show the raw missed detection and false alarm rates. The third column shows the fraction of cases in which the likelihood ratio failed to reach either the alarm or dismissal limit within the short span of measurements that were available. The fourth column shows the fraction of cases in which the test first reached or exceeded one boundary, either alarm or dismissal, and then relaxed back into the in-decision region between the two boundaries. The final column requires a bit more explanation. In practice, if the WSPRT has not terminated with a decision by the time we process the last measurement, we assume that the prudent course of action is to maneuver. In some cases, this would prove to be a correct decision, while in others it would result in a false alarm. We report this increased false alarm rate in the final column.

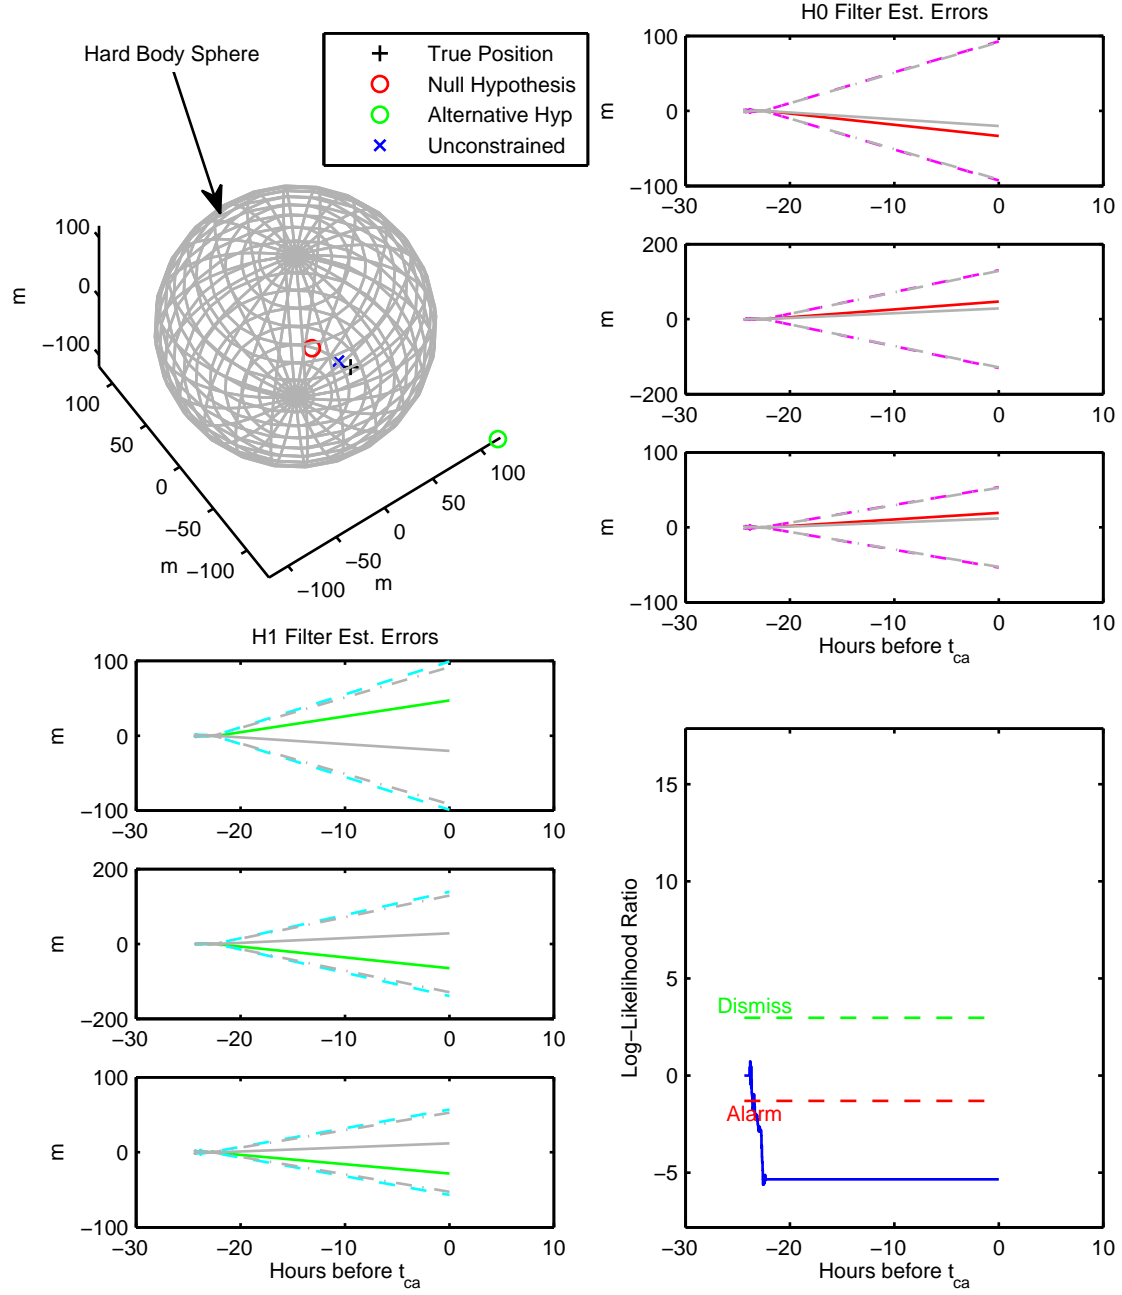


Figure 6. Clear Hit Case ($\|r(t_{ca})\| = 74 \text{ m}$)




Figure 7. Near hit case ($\|r(t_{ca})\| = 114 \text{ m}$)

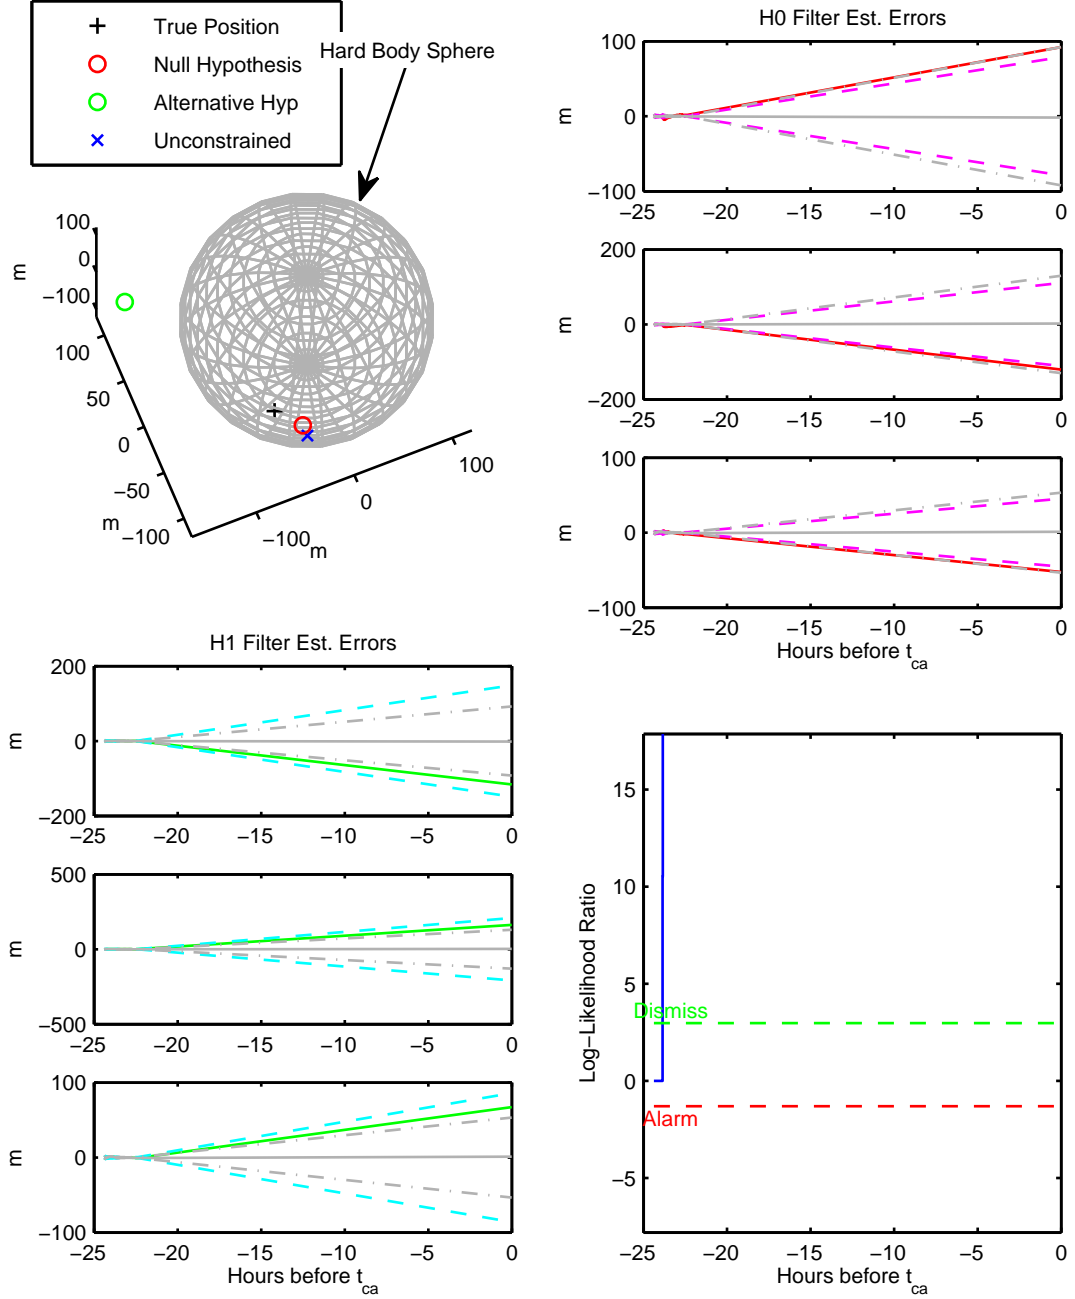


Figure 8. Near miss case ($\|r(t_{ca})\| = 132 \text{ m}$)

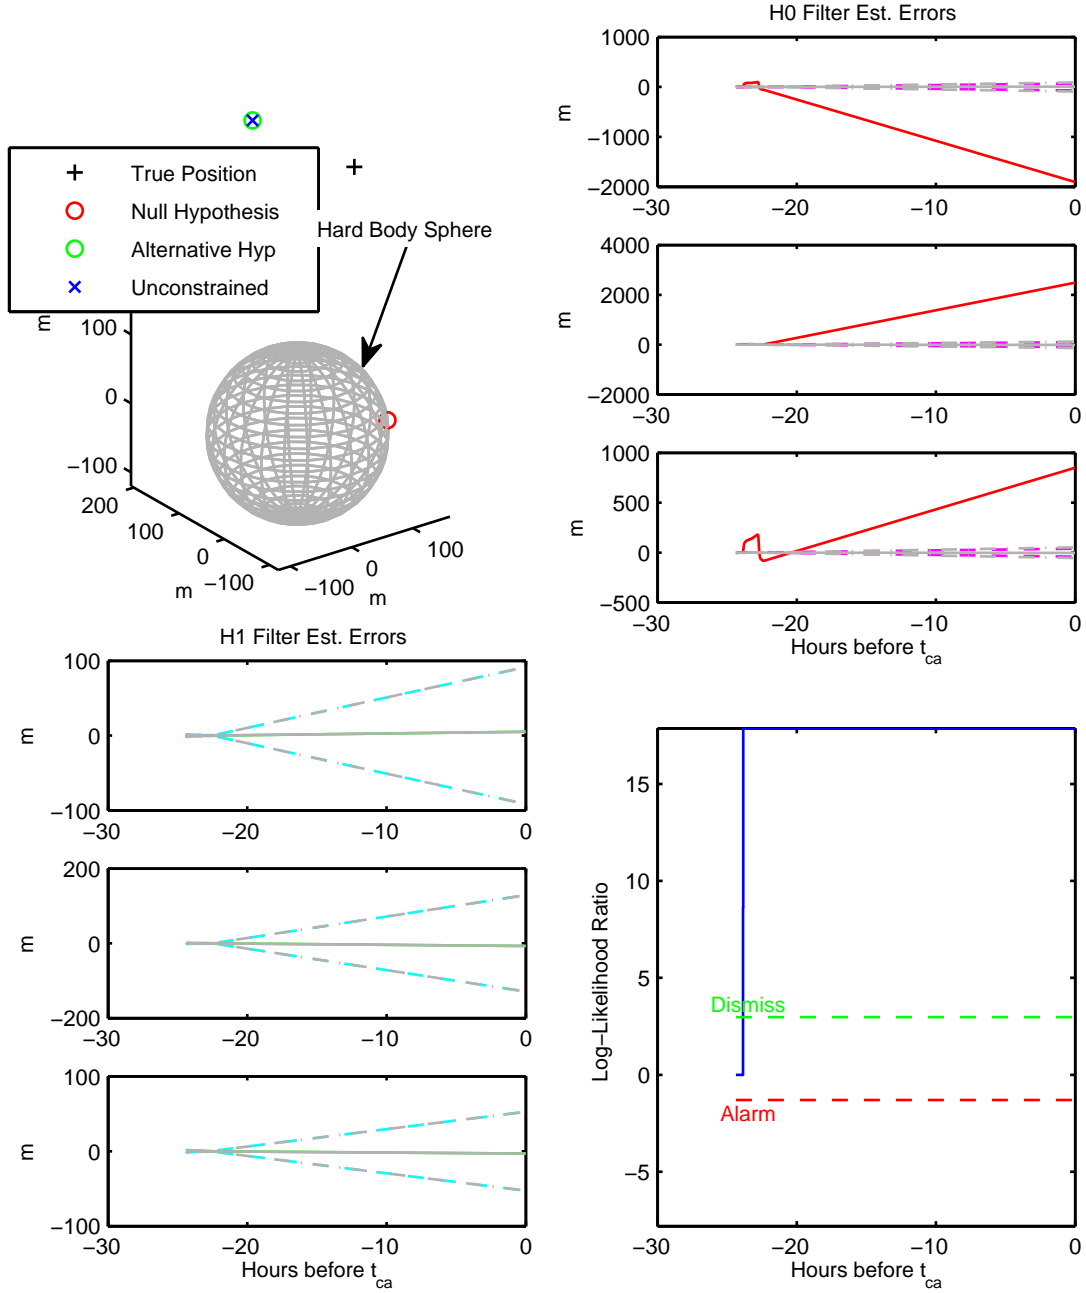


Figure 9. Clear miss case ($\|r(t_{ca})\| = 359 \text{ m}$)

If, per Wald, we assume that any time the WSPRT reaches a decision limit, we would terminate the test, such a policy would remove the possibility of an indecision case occurring. This is the policy we used in constructing the statistics of Table 1, although we nonetheless continued running the test in order to collect the indecision statistics as well. We note that in many cases it may be minimally costly to wait until the end of the data arc to make a recommendation; if the filters are continuing to improve in accuracy as they process more data, the final value of the likelihood ratio is likely to be more accurate and hence both the power and the specificity of the test would improve.

One can see in Table 1 that the missed detection rates slightly exceed the desired rate of 1/1000, while the raw false alarm rates are well within the desired false alarm rate of 1/20. While the elevated effective false alarm rates are not unexpected given our policy of raising the alarm in a no-decision case, the slightly elevated missed detection statistics are notable. For row two of the table, the missed detection rate corresponds to seven of the 6,000 trials which were truthfully inside the combined hard body radius, but which were misidentified as being outside. This is only one in excess of the six that we would expect to occur for a test designed to achieve a missed detection rate of 1/1000. We believe it is reasonable to consider this a non-significant fluctuation. On the other hand, for row one of the table, the missed detection rate corresponds to ten of 6,000 trials; this excess of four events versus the six we expect appears to us to be statistically significant⁶. Therefore, we may conclude from Table 1 that use of the projection step only is superior to using both projection and translation, and that when so doing, the WSPRT meets our expectations for its performance.

6. Conclusions

The Wald sequential probability ratio test based on a bank of inequality-constrained current-state extended Kalman filters, as proposed and demonstrated in this paper, is a promising new approach to conjunction assessment. Although it may require the solution of multiple-revolution Lambert problem for multiple sigma-points to project the inequality constraint from the predicted time of closest approach to the current measurement time, we believe this problem is fundamentally easier than trying to predict a probability density into the future and then compute integrals over it. The MMS mission, due to launch in 2014, plans to utilize our approach as part of its conjunction assessment process.

7. References

- [1] Foster, J. L., Jr. and Estes, H. S. "A Parametric Analysis of Orbital Debris Collision Probability and Maneuver Rate for Space Vehicles." Tech. Rep. JSC-25898, NASA Johnson Space Center, Houston, TX, 1992.
- [2] Akella, M. R. and Alfriend, K. T. "Probability of Collision Between Space Objects." *Journal of Guidance, Control and Dynamics*, Vol. 23, No. 5, pp. 769–772, September–October 2000.

⁶Since missed detections are rare events out of a large population in our simulation, one may reasonably approximate their distribution as Poisson, in which case both the mean and variance of the expected number of missed detections would be six.

- [3] Chan, K. “Improved Analytical Expressions for Computing Spacecraft Collision Probabilities.” “Space Flight Mechanics 2003,” Univelt, 2003.
- [4] Carpenter, J. R. and Markley, F. L. “Sequential Probability Ratio Test for Collision Avoidance Maneuver Decisions.” “Proceedings of the Kyle T. Alfriend Astrodynamics Symposium,” Univelt, 2010.
- [5] Wald, A. Sequential Analysis. Dover Publications, 2004.
- [6] Carpenter, J. R., Markley, F. L., Alfriend, K. T., Wright, C., and Arcido, J. “Sequential Probability Ratio Test for Collision Avoidance Maneuver Decisions Based on a Bank of Norm-Inequality-Constrained Epoch-State Filters.” “Astrodynamics 2011,” Advances in the Astronautical Sciences. Univelt, 2011.
- [7] Brown, R. G. and Hwang, P. Y. Introduction to Random Signals and Applied Kalman Filtering. John Wiley and Sons, Inc., New York, NY, 3rd edn., 1997.
- [8] Zanetti, R., Majji, M., Bishop, R. H., and Mortari, D. “Norm-Constrained Kalman Filtering.” Journal of Guidance, Control, and Dynamics, Vol. 32, No. 5, pp. 1458–1465, September–October 2009.
- [9] Julier, S. J. and LaViola, J. J. “On Kalman Filtering with Nonlinear Equality Constraints.” IEEE Transactions on Signal Processing, Vol. 55, No. 6, pp. 2774–2784, June 2006.
- [10] Simon, D. and Chia, T. L. “Kalman filtering with state equality constraints.” IEEE Transactions on Aerospace and Electronic Systems, Vol. 38, No. 1, pp. 128–136, Jan 2002.
- [11] Nørgaard, M., Poulsen, N. K., and Ravn, O. “New Developments in State Estimation for Nonlinear Estimation Problems.” Automatica, Vol. 36, No. 11, pp. 1627–1638, November 2000.
- [12] Nørgaard, M., Poulsen, N. K., and Ravn, O. “Advances in Derivative-Free State Estimation for Nonlinear Systems.” Tech. Rep. IMM–REP–1998–15, Technical University of Denmark, 2800 Lyngby, Denmark, April 7, 2000.
- [13] Bauer, R. “Uniform Sampling of SO_3 .” J. P. Lynch, editor, “2001 Flight Mechanics Symposium,” NASA/CP–2001–209986, pp. 347–359. NASA Goddard Space Flight Center, June 2001.